determining whether the partition was previously processed;

generating a new minibatchid and a new partitionid for a new partition based upon determining the partition was not previously processed;

determining whether a record was previously processed based upon determining the partition was previously processed; and

processing the record of the partition based upon determining the record was not previously processed.

2. The method of claim **1**, wherein generating a new minibatchid and a new partitionid for a new partition based upon determining the partition was not previously processed further comprises processing the record of the new partition.

3. The method of claim **2** further comprises storing the new minibatchID and the new paritionID in a streams table key column and providing a dummy payload

4. The method of claim **1**, wherein the eventid comprises a monotonically increasing event value.

5. The method of claim **1**, wherein assigning the eventID to the corresponding entityID further comprises storing the entityID in a counters table key column and storing an eventID in a counters table payload column of the corresponding to the entityID.

6. The method of claim **1**, wherein determining whether the partition was previously processed comprises locating a match of the minibatchID and partitionID in the streams table key column to the minibatchID and partitionID of the partition.

7. The method of claim **1**, wherein determining whether a record was previously processed comprises locating a match of the minibatchID, partitionID, and entityID in an events table key column and a set of EventIDs in the events table payload column to the minibatchID, partitionID, entityID, and set of EventIDs for the partition.

8. A distributed stream computing apparatus in non-idempotent output operations, the apparatus comprising:

at least one memory operable to store program instructions;

at least one processor operable to read the stored program instructions; and

according to the stored program instructions, the at least one processor is configured to be operated as:

a driver configured to assign an eventid to a corresponding entityid, to determine a minibatchid and a partitionid for a partition, to determine whether the partition was previously processed, to generate a new minibatchid and a new partitionid for a new partition based upon determining the partition was not previously processed, and to determine whether a record was previously processed based upon determining the partition was previously processed; and

one or more compute nodes configured to process the record of the partition based upon the driver determining the record was not previously processed.

9. The apparatus of claim **8**, wherein the driver being configured to generate a new minibatchid and a new partitionid for a new partition based upon determining the partition was not previously processed, one or more compute nodes are further configured to process the record of the new partition.

10. The apparatus of claim **9**, wherein the driver is further configured to store the new minibatchID and the new paritionID in a streams table key column and to provide a dummy payload

11. The apparatus of claim **8**, wherein the eventid comprises a monotonically increasing event value.

12. The apparatus of claim **8**, wherein the driver, being configured to assign the eventID to the corresponding entityID, is further configured to store the entityID in a counters table key column and to store the eventID in a counters table payload column of the corresponding to the entityID.

13. The apparatus of claim **8**, wherein the driver, being configured to determine whether the partition was previously processed, is further configured to locate a match of the minibatchID and partitionID in the streams table key column to the minibatchID and partitionID of the partition.

14. The apparatus of claim **8**, wherein the driver, being configured to determine whether a record was previously processed, is further configured to locate a match of the minibatchID, partitionID, and entityID in an events table key column and a set of EventIDs in the events table payload column to the minibatchID, partitionID, entityID, and set of EventIDs for the partition.

15. A non-transitory computer readable storage medium, implemented by one or more processors, storing traffic segmentation program for causing a computer to function as:

a driver configured to assign an eventid to a corresponding entityid, wherein the eventid comprises a monotonically increasing event value, to determine a minibatchid and a partitionid for a partition, to determine whether the partition was previously processed, to generate a new minibatchid and a new partitionid for a new partition based upon determining the partition was not previously processed, and to determine whether a record was previously processed based upon determining the partition was previously processed; and

one or more compute nodes configured to process the record of the partition based upon the driver determining the record was not previously processed.

16. The non-transitory computer readable storage medium of claim **15**, wherein the driver being configured to generate a new minibatchid and a new partitionid for a new partition based upon determining the partition was not previously processed, one or more compute nodes are further configured to process the record of the new partition.

17. The non-transitory computer readable storage medium of claim **16**, wherein the driver is further configured to store the new minibatchID and the new paritionID in a streams table key column and to provide a dummy payload

18. The non-transitory computer readable storage medium of claim **15**, wherein the driver, being configured to assign the eventID to the corresponding entityID, is further configured to store the entityID in a counters table key column and to store the eventID in a counters table payload column of the corresponding to the entityID.

19. The non-transitory computer readable storage medium of claim **15**, wherein the driver, being configured to determine whether the partition was previously processed, is further configured to locate a match of the minibatchID and partitionID in the streams table key column to the minibatchID and partitionID of the partition.

20. The non-transitory computer readable storage medium of claim **15**, wherein the driver, being configured to determine whether a record was previously processed, is further